



# Using interviewer random effects to remove selection bias from HIV prevalence estimates

## Citation

McGovern, Mark E, Till Bärnighausen, Joshua A Salomon, and David Canning. 2015. "Using Interviewer Random Effects to Remove Selection Bias from HIV Prevalence Estimates." *BMC Medical Research Methodology* 15 (1): 8. doi:10.1186/1471-2288-15-8.

## Published Version

10.1186/1471-2288-15-8

## Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:23586586>

## Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

## Share Your Story

The Harvard community has made this article openly available.  
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

RESEARCH ARTICLE

Open Access

# Using interviewer random effects to remove selection bias from HIV prevalence estimates

Mark E McGovern<sup>1,2\*</sup>, Till Bärnighausen<sup>2,3</sup>, Joshua A Salomon<sup>2</sup> and David Canning<sup>1,2</sup>

## Abstract

**Background:** Selection bias in HIV prevalence estimates occurs if non-participation in testing is correlated with HIV status. Longitudinal data suggests that individuals who know or suspect they are HIV positive are less likely to participate in testing in HIV surveys, in which case methods to correct for missing data which are based on imputation and observed characteristics will produce biased results.

**Methods:** The identity of the HIV survey interviewer is typically associated with HIV testing participation, but is unlikely to be correlated with HIV status. Interviewer identity can thus be used as a selection variable allowing estimation of Heckman-type selection models. These models produce asymptotically unbiased HIV prevalence estimates, even when non-participation is correlated with unobserved characteristics, such as knowledge of HIV status. We introduce a new random effects method to these selection models which overcomes non-convergence caused by collinearity, small sample bias, and incorrect inference in existing approaches. Our method is easy to implement in standard statistical software, and allows the construction of bootstrapped standard errors which adjust for the fact that the relationship between testing and HIV status is uncertain and needs to be estimated.

**Results:** Using nationally representative data from the Demographic and Health Surveys, we illustrate our approach with new point estimates and confidence intervals (CI) for HIV prevalence among men in Ghana (2003) and Zambia (2007). In Ghana, we find little evidence of selection bias as our selection model gives an HIV prevalence estimate of 1.4% (95% CI 1.2% – 1.6%), compared to 1.6% among those with a valid HIV test. In Zambia, our selection model gives an HIV prevalence estimate of 16.3% (95% CI 11.0% - 18.4%), compared to 12.1% among those with a valid HIV test. Therefore, those who decline to test in Zambia are found to be more likely to be HIV positive.

**Conclusions:** Our approach corrects for selection bias in HIV prevalence estimates, is possible to implement even when HIV prevalence or non-participation is very high or very low, and provides a practical solution to account for both sampling and parameter uncertainty in the estimation of confidence intervals. The wide confidence intervals estimated in an example with high HIV prevalence indicate that it is difficult to correct statistically for the bias that may occur when a large proportion of people refuse to test.

**Keywords:** Selection Bias, HIV Prevalence, Missing Data, Heckman Selection Models, Random Effects Estimation

## Background

Estimates of HIV prevalence from serologic testing in nationally representative household surveys have been considered the “gold standard” in developing countries [1]. However, HIV testing participation rates in these surveys may be low. For instance, in the Demographic

and Health Surveys (DHS), participation in HIV testing ranges from 97% for women in Rwanda in 2005, to 63% for men in Malawi in 2004 and in Zimbabwe in 2005 [2]. In the DHS, respondents are generally asked if they will consent to a blood test after completing a verbal interview [3].<sup>a</sup> Although non-participation can result from either not being contacted for interview, or refusing consent to test for HIV after the interview, the latter reason for non-participation is typically much more common [4]. Therefore, in this paper we focus on correcting for missing data arising from refusal to test. If

\* Correspondence: [mcgovern@hsph.harvard.edu](mailto:mcgovern@hsph.harvard.edu)

<sup>1</sup>Harvard Center for Population and Development Studies, 9 Bow Street, Cambridge, MA 02138, USA

<sup>2</sup>Department of Global Health and Population, Harvard T.H. Chan School of Public Health, Boston, MA 02115, USA

Full list of author information is available at the end of the article

HIV prevalence is different among individuals who test and those who do not test, ignoring missing HIV data may lead to biased estimates of population prevalence. Typical approaches to impute the HIV status of respondents who did not test assume that non-participation is missing at random (MAR), or missing at random conditional on observed covariates. In past studies, this approach has generated estimates of population prevalence that are very similar to estimates based only on the subset of respondents who participated in testing [4,5]. However, the assumption that there are no unobserved variables that are correlated with both HIV status and consent to test may be unrealistic. Longitudinal data show that a person's belief about his or her HIV status may be related to actual status, and may influence the likelihood of consenting to an HIV test [6-9]. If this is the case, the missing data can be described as "non-ignorable", and there will be selection bias in conventional population prevalence estimates based on an incorrect assumption of data missing at random [10]. An additional problem with imputation based approaches is that they are likely to result in confidence intervals which are too narrow because they ignore the uncertainty surrounding the unknown relationship between testing and HIV status, which needs to be estimated. The implicit assumption in these imputation approaches that this correlation is zero with certainty is likely to be violated in many cases.

Several methods have been proposed to account for the selection problem when missing data are non-ignorable, including the use of longitudinal data [7,8,11] and alternative testing procedures [12]; however these approaches are not applicable to existing cross-sectional datasets, such as the DHS. Therefore, an attractive alternative to imputation and other similar methods (for example, inverse-probability weighting), is to use a model that explicitly accounts for the selection process, such as that proposed by Heckman [13]. This approach allows us to obtain asymptotically unbiased and consistent estimates of HIV prevalence, even in the presence of unobserved variables which influence non-participation, such as belief about HIV status [14]. By accounting for both non-participation and HIV status equations explicitly in the model (individuals first chose whether to participate, and it is only conditional on this choice that we observe their HIV status), the key benefit of this method is that we do not require the assumption of data missing at random. In the 2007 Zambia DHS survey, estimated HIV prevalence among men who refused consent to test was found to be 53% using the sample selection model, but only 12% in an imputation model [15]. These results seem plausible when viewed in terms of the longitudinal evidence which indicates that those who refuse to test are substantially more likely to be HIV positive [6-9]. The same method has been used to produce selection-corrected HIV prevalence estimates for other

African countries, as well as at Health and Demographic Surveillance Sites [2,16-18]. These findings provide evidence that selection on unobserved characteristics may be an important source of bias in some existing estimates.

However, there are a number of drawbacks associated with the standard implementation of Heckman-type selection models in the current literature. Because of the way the selection and HIV status equations are typically specified, it is often not possible to use this approach because the models can fail to converge due to collinearity problems. Additionally, the small sample properties of the standard models can be poor. Finally, the standard errors which are typically reported in applications of Heckman-type selection models are too narrow because they do not account for parameter uncertainty in the model estimation. The goal of this paper is to introduce a new method which corrects for each of these limitations. Our method is based on an interviewer random effects model, which can be easily implemented in all surveys, even in the problematic cases where HIV prevalence or selection is either very low (<10%) or very high (>90%). In what follows, we explain the drawbacks associated with the standard Heckman-type model, and outline how our new method corrects for these issues. Finally, we illustrate our approach using household survey data from Zambia and Ghana.

## Methods

### Existing Heckman-type selection models for estimating HIV prevalence

Following the framework adopted by Bärnighausen and colleagues, we model consent to test for HIV for person  $i$  with interviewer  $j$  as the observed outcome arising from a latent variable that can be interpreted as the propensity to consent to testing [15]:

$$\begin{aligned} s_{ij}^* &= x_{ij}'\beta_s + z_j'\gamma_s + u_{ij} \\ s_{ij} &= 1 \text{ if } s_{ij}^* > 0, \quad s_{ij} = 0 \text{ otherwise} \end{aligned} \quad (1)$$

where  $s_{ij}$  is a dummy indicator variable for agreeing to test,  $x_{ij}$  are observed characteristics,  $z_j$  are interviewer effects,  $u_{ij}$  is random error, and  $s_{ij}^*$  is an unobserved latent variable. The equation for the HIV status  $h_{ij}$  of individual  $i$  with interviewer  $j$  is:

$$\begin{aligned} h_{ij}^* &= x_{ij}'\beta_h + \varepsilon_{ij} \\ h_{ij} &= 1 \text{ if } h_{ij}^* > 0, \quad h_{ij} = 0 \text{ otherwise} \\ h_{ij} &\text{ observed only if } s_{ij} = 1 \end{aligned} \quad (2)$$

where  $h_{ij}^*$  is again a latent variable,<sup>b</sup> and  $\varepsilon_{ij}$  is an error term. The measured outcomes, consent to test,  $s_{ij}$ , and HIV status,  $h_{ij}$ , are individual level variables. The main independent variables for both equations,  $x_{ij}$ , include predictors at the individual and household level. These socioeconomic, geographic, and demographic characteristics are derived

from the DHS survey data. In order to provide asymptotically unbiased estimates, Heckman-type selection models additionally require a selection variable which predicts consent to test, but not HIV status [19]. In this case, we use interviewer identity,  $z_j$ , as the selection variable, as it is plausibly unrelated to HIV status. Interviewer allocation is most likely a function of survey design, and therefore should not be related to respondents' characteristics, such as HIV status. Therefore, we include interviewer identity as an additional predictor of consenting to test for HIV. In DHS surveys with nested HIV surveys, once a person has been contacted, the interviewer provides her with information about HIV testing and offers her an HIV test during an informed consent process. The identity of the interviewer who sought consent from that participant is recorded in the dataset as an anonymized code, one for each respondent.<sup>c</sup>

Previous papers have also used the identity of interviewers as a selection variable that is correlated with consent to test but is assumed to be unrelated to HIV status [13,16,17]. For those who are absent and not contacted for interview, day of first contact with the household is an alternative selection variable when survey design allows for follow-up attempts to contact missing participants [13].<sup>d</sup> Better interviewers may have a personality type (for example the ability to show empathy for the interviewee), or relevant experience, which increases their participation rates. This assumption is testable, and interviewer identity is highly correlated with consent in our data. Interviewers are not randomly assigned to participants in DHS surveys, but are assigned to specific regions and tend to be matched to respondents by sex and language. However, sex and language (as well as a rich set of other participant characteristics) are recorded in the DHS datasets, and once they have been controlled for, the interviewer assignment is likely random and uncorrelated with participants' HIV status. Assuming the selection variable is valid (interviewer identity predicts consent to test but not HIV status), Heckman-type selection models will provide asymptotically unbiased estimates, even when the data are not missing at random [14]. Simulation studies have confirmed that this method provides reasonable estimates of HIV prevalence in large samples under the assumption that, conditional on observed participant characteristics, interviewer identity is uncorrelated with HIV status [20].

In the standard approach, which we refer to as a fixed effects model, and which has previously been used to adjust for missing data in some HIV surveys [15], the selection variable (interviewer identity) takes the form of a series of binary indicator variables in the selection equation, which take the value 1 if the individual was interviewed by that interviewer, and 0 otherwise. For interviewers  $j = 1 \dots J$ , there will therefore be  $J - 1$  indicator

dummy variables  $z_1 \dots z_{J-1}$ , with the  $J$ th being the omitted category. This Heckman selection approach is an individual level model because the selection and HIV status equations are measured at the individual level (a binary indicator for whether the respondent consented to test, and a binary indicator for whether that respondent was HIV positive among those who consented to test). However, the interviewer indicator variables (fixed effects) take the same value (1 or 0) for all respondents who were interviewed by the same interviewer. In this sense, the model is hierarchical: respondents are nested within interviewers.

As the dependent variables are binary, we adopt a bivariate probit Heckman-type selection model, which models both consent to test and HIV status equations simultaneously in a maximum likelihood framework [21,22]. The error terms in these equations ( $u_{ij}, \varepsilon_{ij}$ ) are assumed to be drawn from a bivariate normal distribution, each with mean zero, variance 1, and correlation parameter  $\rho = \text{corr}(u_{ij}, \varepsilon_{ij})$ . This approach of modelling consent and HIV status jointly allows us to adjust for dependence between selection into testing and HIV status [13].

The key parameter in the model is  $\rho$ . If  $\rho = 0$ , there is no correlation between testing and HIV status once observed variables have been controlled for, and imputation of the HIV status of those who do not test based on their observed characteristics is possible [4,5]. If  $\rho \neq 0$  however, conditional on observed characteristics testing is correlated with HIV status, and the probability of being HIV positive among those who refuse to test will be different from the prevalence rate among people with similar observable characteristics who do test. The key issue in the model is therefore to find robust estimates of  $\rho$ . Regarding inference, the DHS surveys are generally carried out within fixed strata representing urban and rural areas of each region. Within each stratum, a cluster of households is randomly selected from within a set of possible primary sampling units (PSU), usually defined by a preceding census [23]. All of our models use robust standard errors which adjust for this complex sample design by clustering at the PSU level. We also weight our estimates of HIV prevalence to match the national population using the household level weights provided by the DHS.

#### Limitations of existing approaches

There are three main limitations to the use of Heckman-type selection models which include fixed effects for each interviewer as the selection variables which predict consent to test but not HIV status. First, these fixed effects are often not identified for interviewers who conducted few interviews, or who have only successes or failures in consent.<sup>e</sup> Fixed effects may also not be identified for interviewers who only interview people of a particular language

in a particular region, as this can result in their fixed effects ( $z_j$ ) being collinear with these interviewee characteristics ( $x_{ij}$ ). These problematic interviewer fixed effects are typically pooled to a common value by creating a new category for these interviewers, which also includes all interviewers who complete less than 50 interviews. The assumption of a common value for the interviewer effect on consent in these cases is difficult to justify. In addition, even when interviewer fixed effects are formally identified, this identification may be weak and lead to lack of convergence in model estimates.

Second, while maximum likelihood estimators are asymptotically unbiased and consistent, their small sample properties can be poor [24,25]. In a simulation study, the finite sample bias for the maximum likelihood estimate in a related recursive bivariate probit model was particularly large when the true probability of the outcome was near the boundary of zero or one [26]. This bias occurs because the maximum likelihood estimator selects the most likely parameter value (the mode of the likelihood function), and not the expected value of the parameter. When the likelihood function is uni-modal and symmetrical, the maximum likelihood estimate is usually unbiased, at least asymptotically. This case will arise, for example, when estimating regression parameters in the linear model with Gaussian errors where maximum likelihood estimates are identical to ordinary least squares estimates. However, when the parameter space is bounded, and the true parameter is near the boundary (in our case,  $\rho = \pm 1$ ), the likelihood function is usually highly skewed and its mode and mean can be very different. Rates of non-convergence in bivariate probit models can be high when this occurs [27], and the maximum likelihood estimate can be highly implausible in these cases (for example, predicting that everyone who declines to test is HIV positive or HIV negative).<sup>f</sup> This problem is clearly illustrated in our analysis of the Ghanaian data where HIV prevalence is low.

Finally, the analytic standard errors for the HIV prevalence estimates which are typically reported in applications of Heckman-type selection models are too narrow because they do not account for regression parameter uncertainty in constructing the joint model of selection and HIV status. Bootstrap standard errors have been shown to provide reliable inference in the context of bivariate probit models, particularly when the number of observations is low, and either the probability of the outcome occurring or the probability of selection is high or low [26]. Unfortunately, the bootstrap procedure is difficult to implement with the fixed effects approach as the number of interviewees per interviewer will vary with each bootstrap sample, and different sets of interviewer fixed effects become unidentified in each iteration. Alternative parametric simulation approaches require strong assumptions on the correlation

between unobserved error terms [2]. In addition, when the maximum likelihood estimate of  $\rho$  is at the boundary of the parameter space, the assumptions required for asymptotic normality of the maximum likelihood estimator are no longer met, and inference is non-standard. In particular, the usual properties of the bootstrap approach fail [28,29].

#### **New interviewer random effects methodology**

We account for these three issues by proposing a random effects approach with a bias correction for cases which are difficult to estimate due to low or high rates of non-participation or HIV prevalence. Instead of introducing a fixed effect for each interviewer into the selection equation, we assume that interviewer effectiveness is a random effect. Where data have a hierarchical structure, researchers often face the choice between a model based on random effects and a model based on fixed effects. Our approach for estimating HIV prevalence can be viewed in terms of this framework. The fixed effects approach requires us to estimate a separate parameter for each interviewer (a series of binary indicator variables), some of which may not be identified. While the random effects approach has the usual cost of requiring the assumption that the interviewer effects are drawn from some distribution (typically a normal distribution), this model has the advantage that fewer parameters are required, and it is therefore more parsimonious (and potentially more efficient as a result). The random effects approach also generally requires the assumption that the random effects are uncorrelated with the individual level covariates, but this assumption can be relaxed by including the average characteristics of each interviewer's interviewees in the model [30,31].

The random effects approach allows us to estimate interviewer effects for all interviewers, even those which are not identified in the fixed effects method. Using a regression model, we can predict the interviewer's position in the distribution of success in eliciting consent to an HIV test, and include this information in the selection equation as a single continuous variable. There are two approaches to estimating these random effects. The first is to use a hierarchical probit for consent with interviewers at the group level. The random effects can then be estimated directly from this model. The alternative is to calculate the random effects as the average of the error terms for each interviewer in a probit model for consent. We have found almost identical results from both methods, however we prefer the latter approach for practical reasons because it is much less time consuming.

Our procedure is therefore as follows. We first estimate the interviewer random effects from the selection equation only (stage one), then include these constructed parameters as our exclusion restriction in a Heckman-type model

(stage two). We also include the average characteristics of an interviewer's interviewees in the selection equation to capture potential differences in interviewee attributes across interviewers [30,31]. Because our random effects are a single continuous variable, we can estimate the interviewer parameters using this method in cases in which it is difficult to identify and estimate fixed effects. As a consequence, we avoid arbitrary pooling of interviewer effects and allow straightforward bootstrapping to obtain corrected confidence intervals for our HIV prevalence estimates.

We have tested whether including the mean characteristics of the interviewer's interviewees are required in the model to account for potential correlation between interviewers and the average characteristics of their interviewees. We regress whether a respondent consented to test for HIV on the control variables used in the model ( $x_{ij}$ ), the interviewer fixed effects ( $z_j$ , binary indicator variables for whether the respondent was interviewed by that interviewer), and the average values of the control variables for that interviewer ( $\bar{x}_j$ ). We then test whether the coefficients on  $\bar{x}_j$  are jointly equal to 0. For both Zambia and Ghana, our tests reject this null hypothesis, indicating that the  $\bar{x}_j$  are necessary to avoid bias in the estimation [32]. Therefore, we include them in the model for Zambia and Ghana. However, there may be other contexts where the  $\bar{x}_j$  are not necessary, and excluding them could be more efficient. We therefore we suggest a two-stage process: first, testing the relevance of the mean interviewee characteristics for consent; second, including them in the selection model along with the interviewer random effect, but only if they are in fact relevant for consent. If not, the interviewer random effects could be used as the selection variable without this correction.


Our cluster bootstrap procedure is performed by resampling at the level of the primary sampling unit (PSU), reflecting the two stage sampling design adopted by the DHS [23]. As respondents are nested within PSU clusters, it is important to account for this dependence when estimating variance parameters [33]. Therefore, with each bootstrap replication we re-calculate the interviewer random effects. However, there are alternative approaches to bootstrapping in hierarchical models [34], and, as we describe above, it is also possible to conceptualize respondents as being nested within interviewers in our model. Therefore, we have also estimated cluster bootstrap models where we resample at the interviewer level. We find very similar confidence intervals in this analysis.

In addition to specifying the interviewer parameters as a random effect, we also propose a bias correction procedure which improves on the maximum likelihood estimator by including information on all values of the correlation


between the error terms in testing and HIV status which have non-zero probability. In contrast, the maximum likelihood estimate (by definition, the value with the highest posterior probability) ignores less likely values, even where their posterior probability is positive. To implement this procedure, we calculate the likelihood for each value of  $\rho$  from  $-1$  to  $+1$ . Given a uniform (or flat) prior probability distribution, an appropriate transformation results in an approximation to a posterior density function, which allows us to calculate the expected value. We use this uniform prior probability distribution for the following reasons. First, there is little evidence to justify other prior probability distributions. In this context we lack information on whether a particular set of values for  $\rho$  (the correlation parameter of interest) should be preferred, and the use of a uniform prior probability distribution is common [35]. We have little evidence to expect any particular value for  $\rho$  to be more realistic than another. Second, from a practical perspective, the use of a uniform prior probability distribution makes it straightforward to implement our approach in standard software because we can obtain the posterior density function directly from the likelihood. We have also considered a fully Bayesian model for all parameters; however this approach would require additional prior probability distributions for all parameters and the implementation of a simulation procedure such as Markov Chain Monte Carlo [36]. Our methodology is designed to be easy to implement so that researchers can produce both prevalence estimates and valid confidence intervals given their survey data in a straightforward way. Finally, we note that as the sample size becomes large, we expect the choice of prior probability distribution to have a negligible effect on the resulting estimates [37]. However, as more reliable evidence becomes available on the relationship between testing and HIV status in different contexts, future research could consider adopting alternative approaches to specifying the prior probability distribution.

This bias correction approach avoids the problem of boundary estimates, and is feasible to implement even when the likelihood function is monotonic. The difficulties associated with the maximum likelihood estimate, and the rationale for the bias correction, are illustrated in Figures 1 and 2, which demonstrate that when the likelihood is skewed, the most likely value for the correlation parameter and the expected value can differ substantially.

A major source of convergence failure in Heckman-type selection models (once the convergence problems caused by interviewer identity have been solved by the random effects approach) is in estimating the correlation parameter,  $\rho$ . The advantage of the bias correction approach is that we fix  $\rho$  at a particular value, and then estimate the corresponding log likelihood, so  $\rho$  does not have to be estimated in the model. Therefore, it should be possible to implement



**Figure 1** Posterior probability distribution for the correlation between HIV Testing and HIV Status in Zambia 2007 (Men). Graph shows the posterior probability distribution for the correlation between testing and HIV status  $\rho = corr(u_{ij}, \varepsilon_{ij})$ , calculated using a selection model with a flat prior probability distribution over the interval  $[-1,1]$ , and interviewer random effects as the exclusion restriction. The standard maximum likelihood (ML) estimate is shown, as well as the bias corrected estimate which is the mean of the posterior probability distribution. Also shown is the 95% bootstrap confidence interval for the bias corrected estimate, based on 1,000 replications. The bootstrap confidence interval is calculated using the empirical distribution of bootstrap estimates. Details of the statistical procedure are outlined in the appendix (see Additional file 1). Source: DHS Zambia 2007 (men).



**Figure 2** Posterior probability distribution for the correlation between HIV Testing and HIV Status in Ghana 2003 (Men). Graph shows the posterior probability distribution for the correlation between testing and HIV status  $\rho = corr(u_{ij}, \varepsilon_{ij})$ , calculated using a selection model with a flat prior probability distribution over the interval  $[-1,1]$ , and interviewer random effects as the exclusion restriction. The standard maximum likelihood (ML) estimate is shown, as well as the bias corrected estimate which is the mean of the posterior probability distribution. Also shown is the 95% bootstrap confidence interval for the bias corrected estimate, based on 1,000 replications. The bootstrap confidence interval is calculated using the empirical distribution of bootstrap estimates. Details of the statistical procedure are outlined in the appendix (see Additional file 1). The fact that the maximum likelihood estimate lies outside the bootstrap confidence interval for the bias corrected estimate reflects the fact that the posterior distribution has a long left hand tail which is not accounted for by the standard maximum likelihood estimator, and that we use the empirical distribution of the bootstrap estimates to allow for asymmetry when calculating the confidence interval. Source: DHS Ghana 2003 (men).

this method even when the fixed effects and random effects models fail. Full details of the statistical approach are provided in the appendix (see Additional file 1).

### Data

As the main goal of this paper is to introduce a new method for estimating HIV prevalence, we focus on two datasets from the Demographic and Health Surveys, Zambia (2007) and Ghana (2003), which include nationally representative HIV surveys of men aged 15–59. These data were chosen for two reasons. The first reason is that the results of our random effects model can be compared with the previous fixed effects approach which has already been applied to Zambia. The second reason is that the low HIV prevalence in Ghana makes estimation potentially problematic. Using the standard fixed effects approach, we find the correlation between being HIV positive and testing in Ghana to be large and positive (close to a correlation of  $\rho = 1$ ), so that the maximum likelihood estimate for HIV prevalence among those who refuse to test is close to zero. In addition, in scenarios where the correlation between HIV status and non-participation in testing is large in absolute value, the usual models tend to fail to converge [20,26,27], making implementation of standard selection models difficult.

The analysis samples consist of 7,134 men in Zambia, and 5,334 men in Ghana. Table 1 shows the numbers testing for HIV, the numbers who undertook a DHS interview but refused consent to test for HIV, and the numbers who did not undertake an individual interview and did not test (non-contacts). Overall, participation rates in HIV testing were 72% in Zambia and 80% in Ghana. In this paper, we focus on those who refuse to consent to test. In principle, rather than having two separate models – one for contact and one for consent conditional on contact, we should have a single sequential model since requesting a test only occurs after people are interviewed [18]. However, non-contact is relatively rare, and previous studies found little evidence of selection bias among those who were not interviewed in Zambia [15]. We therefore impute HIV prevalence for non-contacts based on the characteristics of the person

as reported by the household respondent, and the assumption that their HIV status is missing at random conditional on these covariates.

A summary of the datasets used in the analysis are available from the DHS for both Zambia (<http://dhsprogram.com/pubs/pdf/FR211/FR211%5Brevised-05-12-2009%5D.pdf>) and Ghana (<http://dhsprogram.com/pubs/pdf/FR152/FR152.pdf>). We follow the approach adopted in previous papers used to prepare the DHS data for analysis of missing data in HIV surveys, the Stata code for which is publicly available from <http://hdl.handle.net/1902.1/17657> [2,15]. Table A1 in the appendix lists the variables used in the analysis and a description of their source in the data (see Additional file 1). Table A2 in the appendix presents information on the median number of interviewees, median consent rates, and median HIV prevalence rates by interviewer. In Zambia, 89 interviewers attempted to obtain consent to an HIV test. The median interviewer conducted 25 interviews, obtained a consent rate of 84%, and of those who consented to test for the median interviewer, 10% were HIV positive. In Ghana, 55 interviewers attempted to obtain consent to an HIV test. The median interviewer conducted 114 interviews, obtained a consent rate of 87%, and of those who consented to test for the median interviewer, 1% were HIV positive. A potential limitation of the random effects model in general is that when there are small number of groups, non-convergence can occur. However, in our data (and in many similar ones, such as most DHS), we have a relatively large number of interviewers (89 in Zambia and 55 in Ghana).

Interviewers who obtain 100% consent rates may be particularly informative, because there is no selection bias for these interviewers. Unfortunately, in our datasets (Zambia and Ghana), very few interviewers obtain 100% consent rates. Likewise, there are almost no interviewers with 0% consent rates. Of the 89 interviewers in total who attempted to obtain consent in Zambia, there were 21 interviewers with 100% consent rates and one with 0% consent; however most of these interviewers conducted a single interview, and all of them conducted less than 10 interviews each. Less than 2% of respondents were interviewed by one of these interviewers. Similarly, for Ghana relatively few respondents were interviewed by an interviewer who obtained 100% or 0% consent rates. These distributions highlight the major advantage of our random effects model, which makes dealing with these interviewers with few interviews much easier from a technical perspective than in the fixed effects framework: we are able to estimate an interviewer effect for all interviewers, even for those who conducted only a single interview. However, because of their limited number, it is difficult to obtain much information from these interviewers in our data, and they are not very informative

**Table 1 Sample size for HIV prevalence estimation among Men in Zambia (2007) and Ghana (2003)**

	Zambia		Ghana	
	N	%	N	%
Observed HIV Status	5,163	72%	4,271	80%
Missing HIV Status (Consent Refused)	1,318	18%	743	14%
Missing HIV Status (No Contact)	653	9%	320	6%
Total	7,134	100%	5,334	100%

Source: Demographic and Health Surveys.



for estimating HIV prevalence in this case. Further details on the interviewers are provided in Figures A1 – A4 in the appendix (see Additional file 1).

### Ethics approval

This study obtained an exemption from full ethics committee review at Harvard University. All data are publicly available and were analyzed anonymously.

### Results and discussion

We estimate HIV prevalence using three different selection models for Zambia and Ghana. First, we replicate the standard interviewer fixed effects method. Second, we implement the interviewer random effects approach without bias correction, and third, we use interviewer random effects with the bias correction. In addition, we compare these results to an imputation based model and analysis of cases with a valid HIV test (observations with complete data, where we exclude respondents with missing values).

Table 2 shows results for Zambia. The imputation model generates an estimate of 12% for population HIV prevalence, which is substantially smaller than the 20% estimate from the fixed effects model. The random effects and bias correction models give estimates that fall in between these two estimates (16%). Comparing the bootstrap and analytic standard errors highlights how the precision of these estimates is greatly overstated when not accounting for uncertainty in the estimation of  $\rho$ . Figure 1 shows the posterior probability distribution for the correlation parameter in Zambia, along with the maximum likelihood estimate, the bias corrected estimate based on the mean of the posterior probability distribution, and associated bootstrap confidence intervals. In this case, the posterior probability distribution is close to being symmetric, and the mode (the maximum likelihood estimate) is close to the mean.

In Ghana, the maximum likelihood estimator is difficult to implement because of low HIV prevalence. In addition,

the relatively small sample size and low HIV prevalence may induce bias. Figure 2 illustrates the estimated posterior probability distribution for  $\rho$ . The concentrated likelihood function is approximately monotonic, resulting in maximum likelihood estimates which are close to the boundary of the parameter space,  $\rho = 1$ , and the implausible prediction that almost all those who decline to test are HIV negative. The maximum likelihood estimate places zero weight on all the values in the left hand tail of the posterior distribution in Figure 2 as they are less probable than the most likely value. To obtain our bias corrected estimate, we integrate over the posterior probability density function to account for all values of  $\rho$  which have positive probability. Under the assumption that the prior probability distribution is correct, this will result in an unbiased estimate for  $\rho$ . We obtain a value of around 0.6, indicating that individuals with HIV are more likely to consent to test in Ghana.

Table 3 reports our population prevalence estimates for Ghana. Relatively narrow confidence intervals reflect the low HIV prevalence in Ghana, and the lack of potential variation in HIV status in bootstrap replicate samples. For the selection models, we find HIV prevalence of 1.4%, compared to 1.6% in the case of the imputation model. Our population prevalence estimates are therefore slightly lower than those obtained from imputation methods. The positive selection we find in Ghana has also been found in other contexts, although it is in contrast to the negative selection found in Zambia [15]. Differences in selection mechanisms may reflect differences in HIV prevalence and other country-specific factors. For example, being HIV positive is unlikely to be a common reason for non-consent in a country with very low HIV prevalence. However, the differences in prevalence estimates are small in absolute terms, and therefore we conclude that there is little evidence that prevalence estimates are affected by selection bias among Ghanaian men. Tables showing parameter

**Table 2 Estimates of HIV prevalence among Men in Zambia (2007)**

Model	HIV prevalence	Analytic 95% CI		Bootstrap 95% CI	
All Men - Fixed Effects Selection Model	20.1%	19.0%	21.3%		
All Men - Random Effects Selection Model	16.3%	15.3%	17.3%	11.0%	18.4%
All Men - Random Effects Bias Correction Selection Model	15.5%	14.5%	16.5%	10.2%	17.9%
Men with Valid HIV Tests	12.1%	11.0%	13.3%		
Men with No Contact - Imputation Model	15.3%	14.2%	16.3%		
All Men - Imputation Model	12.3%	11.4%	13.2%		

In the Heckman-type selection models (rows 1-3), consent to test and HIV status are jointly estimated using a bivariate probit with the following covariates: education, household wealth quintile, type of location, marital status, had a sexually transmitted disease, age at first intercourse, had high risk sex, number of partners, condom use, would care for an HIV-infected relative, knows someone who died of AIDS, previously tested for HIV, smokes, drinks alcohol, language, age group, region, ethnicity and religion. The selection variable which predicts consent but not HIV status is interviewer identity. Full parameter estimates are presented in tables A4-A6 in the appendix (Additional file 1). Analytic standard errors are shown for the fixed effects and random effects models, with bootstrap errors for random effects and random effects bias correction models based on 1,000 replications. Our cluster bootstrap takes account of survey design by drawing a fixed number of clusters (the same as in the original data) from each stratum in each sample. Results from an imputation model are also shown in rows 5–6, along with estimates only using those without missing data (respondents with a valid HIV test). HIV prevalence estimates are weighted. Source: DHS Zambia 2007 (men).

**Table 3 Estimates of HIV Prevalence among Men in Ghana (2003)**

Model	HIV prevalence	Analytic 95% CI		Bootstrap 95% CI	
All Men - Fixed Effects Selection Model	1.4%	1.1%	1.7%		
All Men - Random Effects Selection Model	1.4%	1.1%	1.7%	1.2%	1.6%
All Men - Random Effects Bias Correction Selection Model	1.4%	1.1%	1.7%	1.2%	1.6%
Men with Non-Missing Data (Valid HIV Test)	1.6%	1.2%	2.0%		
Men with No Contact - Imputation Model	1.6%	1.4%	1.8%		
All Men - Imputation Model	1.6%	1.3%	2.0%		

Consent to test and HIV status are jointly estimated using a bivariate probit with the following covariates: education, wealth quintile, marital status, had a sexually transmitted disease, age at first intercourse, had high risk sex, number of partners, condom use, would care for an HIV-infected relative, knows someone who died of AIDS, previously tested for HIV, smokes, language, age group, region, ethnicity and religion. The selection variable which predicts consent but not HIV status is interviewer identity. Full parameter estimates are presented in tables A8-A10 in the appendix (see Additional file 1). Analytic standard errors are shown for the fixed effects and random effects models, with bootstrap errors for random effects and random effects bias correction models based on 1,000 replications. Our cluster bootstrap takes account of survey design by drawing a fixed number of clusters (the same as in the original data) from each stratum in each sample. Results from an imputation model are also shown in rows 5–6, along with estimates only using those without missing data (respondents with a valid HIV test). HIV prevalence estimates are weighted. Source: DHS Ghana 2003 (men).

estimates from all models are presented in the appendix (see Additional file 1).

### Conclusions

This paper confirms that non-participation in HIV testing may be an important source of bias in HIV prevalence estimation that does not correct for non-ignorable missing data. We introduce a random effects approach for Heckman-type selection models, which improves on previous fixed effects approaches by allowing us to estimate interviewer effects for all interviewers, even for interviewers with few interviewees. This approach makes it possible to use bootstrapping to calculate confidence intervals that account for the fact that the relationship between selection and HIV status is uncertain. We also introduce a bias correction model which facilitates estimation of the correlation between consent to test and HIV status when HIV prevalence or non-participation is very high or very low, and when the usual maximum likelihood model fails to converge. In Zambia, we find that men with HIV are less likely to consent to an HIV test. For Ghana, we find little evidence that conventional methods understate HIV prevalence.

An important result from our empirical analysis is that the corrected confidence intervals around the HIV prevalence point estimate can be very wide. These wide confidence intervals accurately reflect the fact that it is difficult to correct statistically for the bias that may occur when many people refuse to test for HIV. As long as consent rates are low, uncertainty in HIV prevalence estimation will likely remain high. It is important not to understate this uncertainty, and our approach provides a practical solution to account for both sampling and parameter uncertainty in the estimation of HIV prevalence confidence intervals when using Heckman-type selection models to remove selection bias.

The goal of this paper is to address key limitations of standard Heckman-type selection models when a valid

selection variable – in this case, interviewer identity – is available. While it is plausible that interviewer allocation is only affected by survey design and not associated with respondents' HIV status, this claim is not possible to prove conclusively without additional data, because the HIV status of those who do not participate in testing is not observed. As the resulting HIV prevalence prediction relates to a population for whom we generally never get to observe true HIV status, it is important to independently validate the model. We are therefore working to obtain objective data in the form of mortality or antiretroviral treatment records with which we can do so [38].

A drawback of the random effects procedure is that it is performed in two steps: first, the interviewer random effects are estimated from a probit model, and then the interviewer random effects are included in the selection model as the selection variable. This additional step of requiring the interviewer effects to be estimated could introduce measurement error into the model. This may be reflected in the slightly lower HIV prevalence estimates we find for the random effects model in Zambia compared to the fixed effects model. Developing the methodology to incorporate the interviewer random effects directly into bivariate probit selection models, or other approaches to eliminate this potential attenuation bias, would be a useful advance. Another general drawback of selection models for binary outcomes is that they require a parametric assumption. We are pursuing alternative methods in order to determine the robustness of these results to violation of the assumption that the error terms in the selection and HIV status equations are distributed as bivariate normal. These limitations, coupled with the wide confidence intervals we find in the empirical analyses, suggests that despite methodological improvements to remove selection biases from HIV prevalence estimates, it is critical to increase consent rates in HIV surveys. In countries with high non-participation in HIV surveys, interventions to increase consent, such as

financial incentives, should be considered for routine implementation.

## Endnotes

<sup>a</sup>Further details of the testing procedure are available online at <http://dhsprogram.com/pubs/pdf/OD61/OD61.pdf>.

<sup>b</sup>Which can be thought of as reflecting propensity to be infected.

<sup>c</sup>The identity of the interviewer who sought to contact eligible participants is also typically available in the DHS as an anonymized code.

<sup>d</sup>Therefore, although we mainly refer to declining consent to test for HIV in this paper, it is straightforward to apply this approach to missing data which arises through other mechanisms which result in non-participation.

<sup>e</sup>Because there is no within-interviewer variation in consent for these interviewers, and these interviewer parameters perfectly predict consent to test for their interviewees.

<sup>f</sup>Practical solutions to the boundary problem in bivariate probit models have been proposed, but these do not necessarily solve the difficulty of being left with implausibly large correlation coefficients [27,39].

## Additional file

**Additional file 1: Further description of the statistical approach and additional results.**

### Abbreviations

HIV: Human immunodeficiency virus; DHS: Demographic and Health Surveys; MAR: Missing at random; ML: Maximum likelihood; CI: Confidence interval; PSU: Primary sampling unit.

### Competing interests

The authors declare that they have no competing interest.

### Authors' contributions

MM, TB, JS, and DC contributed to the design of the paper, and the writing and revision of the manuscript. MM prepared and analyzed the data. All authors read and approved the final manuscript.

### Acknowledgements

We are grateful to Bart Hamilton, two referees, and seminar participants at Harvard, Oxford, PAA 2013, ISNE 2013, the 2013 Health Econometrics Workshop, and the UNAIDS Reference Group on Estimates, Modelling and Projections for comments and suggestions. MM was funded by the Program on the Global Demography of Aging, which receives funding from the National Institute on Aging, Grant No. 1 P30 AG024409-09. TB was funded by the Wellcome Trust. The funding bodies had no role in the design, collection, analysis, and interpretation of data; or in the writing of the manuscript; or in the decision to submit the manuscript for publication.

### Author details

<sup>1</sup>Harvard Center for Population and Development Studies, 9 Bow Street, Cambridge, MA 02138, USA. <sup>2</sup>Department of Global Health and Population, Harvard T.H. Chan School of Public Health, Boston, MA 02115, USA.

<sup>3</sup>Wellcome Trust Africa Centre for Health and Population Studies, University of KwaZulu-Natal, Mtubatuba, South Africa.

Received: 9 July 2014 Accepted: 6 January 2015

Published: 5 February 2015

## References

1. Boerma JT, Ghys PD, Walker N. Estimates of HIV-1 prevalence from national population-based surveys as a new gold standard. *Lancet*. 2003;362:1929–31.
2. Hogan DR, Salomon JA, Canning D, Hammit JK, Zaslavsky AM, Bärnighausen T. National HIV prevalence estimates for sub-Saharan Africa: controlling selection bias with Heckman-type selection models. *Sex Transm Infect*. 2012;88 Suppl 2:i17–23.
3. Mishra V, Vaessen M, Boerma J, Arnold F, Way A, Barrere B, et al. HIV testing in national population-based surveys: experience from the Demographic and Health Surveys. *Bull World Health Organ*. 2006;84:537–45.
4. Marston M, Harris K, Slaymaker E. Non-response bias in estimates of HIV prevalence due to the mobility of absentees in national population-based surveys: a study of nine national surveys. *Sex Transm Infect*. 2008;84 Suppl 1:i71–7.
5. Mishra V, Barrere B, Hong R, Khan S. Evaluation of bias in HIV seroprevalence estimates from national household surveys. *Sex Transm Infect*. 2008;84 Suppl 1:i63–70.
6. Bärnighausen T, Tanser F, Malaza A, Herbst K, Newell M. HIV status and participation in HIV surveillance in the era of antiretroviral treatment: a study of linked population-based and clinical data in rural South Africa. *Trop Med Int Health*. 2012;17:e103–10.
7. Reniers G, Eaton J. Refusal bias in HIV prevalence estimates from nationally representative seroprevalence surveys. *AIDS*. 2009;23:621–9.
8. Floyd S, Molesworth A, Dube A, Crampin AC, Houben R, Chihana M, et al. Underestimation of HIV prevalence in surveys when some people already know their status, and ways to reduce the bias. *AIDS*. 2013;27:233–42.
9. Obare F. Nonresponse in repeat population-based voluntary counseling and testing for HIV in rural Malawi. *Demography*. 2010;47:651–65.
10. Conniffe D, O'Neill D. Efficient Probit Estimation with Partially Missing Covariates. *Adv Econ*. 2011;27:209–45.
11. Arpino B, Cao ED, Peracchi F. Using panel data for partial identification of human immunodeficiency virus prevalence when infection status is missing not at random. *J R Stat Soc Ser A*. 2014;177:587–606.
12. Hund L, Pagano M. Estimating HIV prevalence from surveys with low individual consent rates: annealing individual and pooled samples. *Emerg Themes Epidemiol*. 2013;10:2.
13. Heckman JJ. Sample selection bias as a specification error. *Econometrica*. 1979;47:153–61.
14. Vella F. Estimating models with sample selection bias: a survey. *J Hum Resour*. 1998;33:127–69.
15. Bärnighausen T, Bor J, Wandira-Kazibwe S, Canning D. Correcting HIV prevalence estimates for survey nonparticipation using Heckman-type selection models. *Epidemiology*. 2011;22:27–35.
16. Janssens W, van der Gaag J, de Wit TFR, Tanović Z. Refusal bias in the estimation of HIV prevalence. *Demography*. 2014;51:1131–57.
17. Reniers G, Araya T, Berhane Y, Davey G, Sanders EJ. Implications of the HIV testing protocol for refusal bias in seroprevalence surveys. *BMC Public Health*. 2009;9:1–9.
18. Clark SJ, Houle B. Validation, replication, and sensitivity testing of Heckman-type selection models to adjust estimates of HIV prevalence. *PLoS One*. 2014;9:e112563.
19. Madden D. Sample selection versus two-part models revisited: the case of female smoking and drinking. *J Health Econ*. 2008;27:300–7.
20. Clark SJ, Houle B. Evaluation of Heckman selection model method for correcting estimates of HIV prevalence from sample surveys via realistic simulation. Center for Statistical Social Science, Working Paper No 120. University of Washington; 2012.
21. Dubin JA, Rivers D. Selection bias in linear regression, logit and probit models. *Sociol Methods Res*. 1989;18:360–90.
22. Van de Ven WP, Van Praag B. The demand for deductibles in private health insurance: a probit model with sample selection. *J Econom*. 1981;17:229–52.
23. Corsi DJ, Neuman M, Finlay JE, Subramanian S. Demographic and health surveys: a profile. *Int J Epidemiol*. 2012;41:1602–13.
24. Greene W. The behaviour of the maximum likelihood estimator of limited dependent variable models in the presence of fixed effects. *Econom J*. 2004;7:98–119.
25. Greenland S. Small-sample bias and corrections for conditional maximum-likelihood odds-ratio estimators. *Biostatistics*. 2000;1:113–22.

26. Chiburis RC, Das J, Lokshin M. A practical comparison of the bivariate probit and linear IV estimators. *Econ Lett.* 2012;117:762–6.
27. Butler JS. Estimating the correlation in censored probit models. *Rev Econ Stat.* 1996;78:356–8.
28. Andrews DW. Estimation when a parameter is on a boundary. *Econometrica.* 1999;67:1341–83.
29. Andrews DW. Inconsistency of the bootstrap when a parameter is on the boundary of the parameter space. *Econometrica.* 2000;68:399–405.
30. Mundlak Y. On the pooling of time series and cross section data. *Econometrica.* 1978;46:69–85.
31. Chamberlain G. Analysis of covariance with qualitative data. *Rev Econ Stud.* 1980;47:225–38.
32. Wooldridge JM. *Econometric analysis of cross section and panel data.* Cambridge, Massachusetts: MIT Press; 2010.
33. Wooldridge JM. Cluster-sample methods in applied econometrics. *Am Econ Rev.* 2003;93:133–8.
34. Van der Leeden R, Meijer E, Busing FM. Resampling multilevel models. In: *Handbook of multilevel analysis.* New York: Springer; 2008. p. 401–33.
35. Qian SS, Stow CA, Borsuk ME. On monte carlo methods for Bayesian inference. *Ecol Model.* 2003;159:269–77.
36. Hall BL, Hamilton BH. New information technology systems and a Bayesian hierarchical bivariate probit model for profiling surgeon quality at a large hospital. *Q Rev Econ Finance.* 2004;44:410–29.
37. O'Hagan A, Forster J, Kendall MG. *Bayesian inference.* London: Arnold; 2004.
38. Nyirenda M, Zaba B, Bärnighausen T, Hosegood V, Newell M-L. Adjusting HIV prevalence for survey non-response using mortality rates: an application of the method using surveillance data from rural South Africa. *PLoS One.* 2010;5:e12370.
39. Terza JV, Tsai W. Censored probit estimation with correlation near the boundary: a useful reparameterization. *Rev Appl Econ.* 2006;2:1–12.

doi:10.1186/1471-2288-15-8

**Cite this article as:** McGovern *et al.*: Using interviewer random effects to remove selection bias from HIV prevalence estimates. *BMC Medical Research Methodology* 2015 15:8.

**Submit your next manuscript to BioMed Central and take full advantage of:**

- Convenient online submission
- Thorough peer review
- No space constraints or color figure charges
- Immediate publication on acceptance
- Inclusion in PubMed, CAS, Scopus and Google Scholar
- Research which is freely available for redistribution

Submit your manuscript at  
[www.biomedcentral.com/submit](http://www.biomedcentral.com/submit)



# Appendix to: Using Interviewer Random Effects to Remove Selection Bias from HIV Prevalence Estimates

Mark E. McGovern<sup>\*†</sup>    Till Bärnighausen<sup>‡</sup>    Joshua A. Salomon<sup>§</sup>    David Canning<sup>¶</sup>

## Additional File 1 - Further description of the statistical approach and additional results

### A1 Statistical Approach

#### A1.1 Standard Heckman-Type Selection Model for Estimating HIV Prevalence

As outlined in the main text, our Heckman-type selection model adopts a bivariate probit approach [1, 2], where we model consent to test for HIV for person  $i$  with interviewer  $j$  as the observed outcome arising from a latent variable that may be interpreted as the propensity to consent to testing [3]:

$$s_{ij}^* = x_{ij}'\beta_s + z_j'\gamma_s + u_{ij} \tag{A1}$$

$$s_{ij} = 1 \text{ if } s_{ij}^* > 0, s_{ij} = 0 \text{ otherwise}$$

where  $s_{ij}$  is a binary indicator for agreeing to test,  $x_{ij}$  are observed characteristics,  $z_j$  are interviewer effects,  $u_{ij}$  is random error, and  $s_{ij}^*$  is an unobserved latent variable.  $\beta_s$  and  $\gamma_s$  are the associated vectors of parameters.

---

\*Corresponding author. Address: 9 Bow Street, Cambridge, MA 02138, USA.

<sup>†</sup>Harvard Center for Population and Development Studies and Department of Global Health and Population, Harvard T.H. Chan School of Public Health. Email: [mcgovern@hsph.harvard.edu](mailto:mcgovern@hsph.harvard.edu).

<sup>‡</sup>Department of Global Health and Population, Harvard T.H. Chan School of Public Health, and Wellcome Trust Africa Centre for Health and Population Studies, South Africa. Email: [tbaernig@hsph.harvard.edu](mailto:tbaernig@hsph.harvard.edu).

<sup>§</sup>Department of Global Health and Population, Harvard T.H. Chan School of Public Health. Email: [jsalomon@hsph.harvard.edu](mailto:jsalomon@hsph.harvard.edu).

<sup>¶</sup>Harvard Center for Population and Development Studies and Department of Global Health and Population, Harvard T.H. Chan School of Public Health. Email: [dcanning@hsph.harvard.edu](mailto:dcanning@hsph.harvard.edu).

The equation for the HIV status of individual  $i$  with interviewer  $j$  is:

$$h_{ij}^* = x_{ij}'\beta_h + \epsilon_{ij} \tag{A2}$$

$$h_{ij} = 1 \text{ if } h_{ij}^* > 0, h_{ij} = 0 \text{ otherwise}$$

$$h_{ij} \text{ observed only if } s_{ij} = 1$$

where  $h_{ij}^*$  is again a latent variable, which can be thought of as reflecting propensity to be infected, and  $\epsilon_{ij}$  is an error term. We assume  $(u_{ij}, \epsilon_{ij})$  are jointly distributed as bivariate normal, each with mean zero, variance 1, and correlation parameter  $\rho = \text{corr}(u_{ij}, \epsilon_{ij})$ .

The inclusion of the interviewer effects,  $z_j$ , which are assumed to affect consent to testing, but not HIV status, is crucial to the model. Without variables in the selection equation that are excluded from the HIV status equation, the model is only identified by non-linearities, and does not provide robust estimates [4].

Even with a suitable exclusion restriction, there are a number of technical problems associated with estimating the model. The first concerns the interviewer effects in  $z_j$ , which typically take the form of a series of binary indicator variables, essentially an interviewer fixed effect [3]. Estimation of these interviewer fixed effects may be difficult because, for interviewers who have only successes or failures in obtaining consent, the parameter is not identified. For example, if an interviewer always obtains consent, we know that the coefficient on the indicator variable for that interviewer in the consent equation is large and positive. However, the coefficient is not identified, as there is no variation in consent for that interviewer. The coefficient for that interviewer could be arbitrarily large, as any very large positive interviewer effect above some threshold will perfectly predict success, and therefore we cannot distinguish between different effect sizes for interviewers who obtain 100% consent. In addition, we control for the region and language of the interviewee. However, in the DHS interviewers usually work in a small number of regions, and an effort is often made to match them with interviewees by language. This may create collinearity between the interviewer effect and the region and language indicator variables in the equation, making it difficult or even impossible to estimate the interviewer effects. The approach adopted by Bärnighausen and colleagues was to limit estimation of individual interviewer effects to those interviewers with more than 50 interviews, and not to use the interviewer effects of those interviewers who had conducted more than 50 interviews when including them lead to identification problems [3].<sup>1</sup> Instead, interviewers with less than 50 interviews, or interviewers whose inclusion caused identification problems, were grouped together as a baseline group with an assumed common interviewer effect.

A conceptual difficulty is that grouping all of the interviewers whose individual interviewer effects are not identified into a baseline group may not be a convincing strategy, because we could be pooling interviewers who always have consent with those who never achieve consent, and assuming that they have the same average effectiveness in obtaining consent. Thus, we do not exploit the fact that the HIV tests made by interviewers who always obtain consent are the most informative

---

<sup>1</sup>The identification problem is usually manifested in practice by the model failing to converge, technically the unidentified parameters have estimated standard errors that are large and increasing at every iteration – the limit is an unbounded standard error on the unidentified parameter.

since they do not suffer from selection bias. As we discuss in the main text, two other limitations with the fixed effect approach are potential small sample bias, and failure of the bootstrap method for calculating standard errors.

## A1.2 Selection Model Using Random Effects

Our approach to these problems is to assume that the unobserved interviewer effects reflect some underlying parametric distribution which describes interviewer effectiveness. If interviewers were randomly assigned to survey households we could simply assume that each interviewer was a random draw from the pool of interviewers. However, systematic matching of interviewers to subjects, particularly by region and language, can mean that there may be a correlation between interviewer success rates and who they are matched with. Therefore, we write the interviewer effect as [5, 6]:

$$z_j = \bar{x}'_j \delta + v_j \tag{A3}$$

$$\bar{x}_j = \sum_i^{n_j} \frac{x_{ij}}{n_j}, v_j \sim N(0, \sigma_z^2)$$

Where  $\bar{x}_j$  represents the average characteristics of the  $n_j$  people that interviewer  $j$  interviews. Excluding these controls could potentially invalidate the exclusion restriction of the model. In order to be asymptotically unbiased and consistent, a requirement of random effects models is that there is no correlation between the random effects,  $v_j$ , and the explanatory variables at the individual level,  $x_{ij}$  [5, 6]. In order for this to hold, we must assume that, controlling for these observable averages, there is no remaining correlation between the error term in equation (A3) and the individual level variables in the model. In particular, we assume that interviewers are not systematically assigned to groups of survey participants based on unobservable characteristics, but only on the observable characteristics measured by the survey. As these  $\bar{x}_j$  likely depend only on survey design, they should not enter the HIV status equation.<sup>2</sup>

This assumption of interviewer random effects gives us the selection equation:

$$s_{ij}^* = x_{ij}' \beta_s + \bar{x}'_j \delta + v_j + u_{ij}$$

$$s_{ij} = 1 \text{ if } s_{ij}^* > 0, s_{ij} = 0 \text{ otherwise} \tag{A4}$$

In principle, we could estimate the system given by equations (A2) and (A4) by maximum likelihood. However, this is difficult as we have a selection equation which has a random effect, requiring numerical integration, inside a bivariate probit model.

There is, however, a simple consistent estimator. Dubin and Rivers show that the bivariate probit model with selection can be estimated by first finding a consistent estimate of the parameters of the selection model, ignoring the covariance of the error terms, and then estimating the parameters

---

<sup>2</sup>If they are included in the HIV status equation they have little effect on our results; they simply reduce the efficiency of estimation.

of the full model by maximum likelihood, holding the selection equation parameters at their first stage estimates [2]. The procedure is as follows. We first estimate the interviewer effects from the selection equation only (stage one), then include these constructed parameters as our exclusion restriction in a Heckman-type selection model (stage two). This two-stage approach is consistent, though not fully efficient, and the second stage does not produce the correct analytic standard errors since the interviewer effects are estimated in the first stage, but treated as exogenous variables in the second stage. Murphy and Topel discuss the consequences for confidence intervals in models which use data that are estimated from a first stage [7].

We can implement this approach by first estimating the interviewer effect as shown in equation (A3). Equation (A4), the selection equation, is then run using the predicted interviewer effect ( $z_j = \bar{x}'_j\delta + v_j$ ) as the exclusion restriction. Assuming random effects avoids the estimation problems of the fixed effects approach; because we are able to estimate an interviewer effect for all interviewers we can implement a bootstrap procedure to adjust our standard errors for the fact that the relationship between consent and HIV status is uncertain and needs to be estimated. The assumption that the interviewer effects are normally distributed random effects around  $\bar{x}'_j\delta$  also means we have a smaller set of parameters to estimate than in the fixed effects model.

Conceptually, we could estimate the interviewer random effects using a multilevel regression for consent, with level one being the individual and level two being the interviewers. However, a simpler method of constructing consistent estimates of the interviewer random effects is to implement a probit model and compute the predicted random effect,  $\hat{v}_j$ , as the average of the error term for each interviewer.<sup>3</sup> We use this probit model to compute the estimated interviewer effect,  $\hat{z}_j = \bar{x}'_j\hat{\delta} + \hat{v}_j$ , where  $\hat{\delta}$  is the vector of estimated coefficients on the interviewer averages from equation (A3), and  $\hat{v}_j$  is the predicted random effect from equation (A3).<sup>4</sup> Having obtained the estimates  $\hat{z}_j$ , we can then estimate the full bivariate probit model in equations (A1) and (A2) using this estimated interviewer effect as the selection variable that affects consent but does not appear in the HIV status equation.

Since the first stage is a consistent estimator of the interviewer effect, the two-stage procedure will be consistent. We can address the problem of incorrect standard errors by bootstrapping over the whole two-stage procedure.

### A1.3 Bias Correction

While the random effects approach solves the identification problems associated with the interviewer parameters and allows bootstrapping, a second problem remains. In small samples, particularly when consent to test is very rare (or very common), or the HIV rate is very low (or very high), it is difficult to estimate the correlation between consent to test and HIV status [8], and the maximum likelihood estimates in the context of bivariate probit models can be biased [9]. Also, the model may fail to converge, or may produce a result of  $\rho = \pm 1$ , on the boundary of the possible parameter space, in which case the assumptions required for asymptotic normality of the maximum likelihood

---

<sup>3</sup>Estimating this first stage as a random effects model is computationally intensive and produces almost identical results to the simple probit. The simple probit produces consistent estimates and this is all that is required for the second stage.

<sup>4</sup>Our results are robust to just using  $\hat{v}_j$  alone as the interviewer effect rather than the full estimate  $\bar{x}'_j\hat{\delta} + \hat{v}_j$



estimator are violated, and standard inference, including bootstrapping, is invalid [10, 11]. Such a result also has the implication that, in terms of predicted probabilities, everyone who fails to test is either HIV positive with certainty ( $\rho = +1$ ), or HIV negative with certainty ( $\rho = -1$ ), which seems implausible. In general, maximum likelihood often does not have desirable finite sample properties [12, 13], as the estimate is the most likely value (in terms of posterior probability), which gives zero weight to values with lower posterior probabilities, even when those probabilities are positive.

We wish to construct an estimate of  $\rho$  that is consistent and also corrects for this small sample bias. Take a data set  $x^*$  and a parameter vector  $\theta^*$ . The likelihood of  $\theta^*$  is simply the probability of the data given these parameters:

$$L(\theta^*) = P(x^*|\theta^*) \tag{A5}$$

In the usual maximum likelihood framework, we are not generally concerned with the likelihood of the observed data given the parameters, we are more interested in the probability of the parameters given the observed data. However, by Bayes rule:

$$P(x^*|\theta^*) = \frac{P(x^*|\theta^*)P(\theta^*)}{P(x^*)} \tag{A6}$$

Where  $P(\theta^*)$  is our prior probability distribution on the parameters. If we have a flat prior probability distribution over the interval  $[-1, 1]$ , so that  $P(\theta^*)$  is the same for every  $\theta^*$ , we have  $P(\theta^*|x^*) \propto P(x^*|\theta^*) = L(\theta^*)$ , and the maximum likelihood estimate of  $\theta^*$  is also the estimate that has the highest posterior probability given the data. While this estimate is the most likely parameter value given the data, we can construct an alternative estimator as:

$$\hat{\theta}^* = E(\theta^*|x^*) = \int_{-\infty}^{+\infty} \theta^* P(\theta^*|x^*) d\theta^* \tag{A7}$$

If the prior probability distribution is flat, we have  $P(\theta^*|x^*) \propto L(\theta^*)$ , the posterior probability is proportional to the likelihood, and we can write our estimator as:

$$\hat{\theta}^* = E(\theta^*|x^*) = \int_{-1}^{+1} \theta^* k L(\theta^*) d\theta^* \tag{A8}$$

Where  $k$  is a normalization factor so that the integral of the likelihood over the parameter space is one, and  $\theta^* k L(\theta^*)$  can be interpreted as an approximation to a probability density function.

The standard maximum likelihood approach chooses the most likely value of  $\theta^*$ , while our approach gives us an average value of  $\theta^*$ , where we average over different models weighted by the probability of the model being correct. This gives consistent estimates (the likelihood function asymptotically puts zero weight on incorrect parameters) and is an unbiased estimator by construction under the assumption that the prior probability distribution is correct. This approach is implemented by calculating the likelihood for each value of  $\rho$ , and then taking the weighted average of  $\rho$ , where the weights are the likelihood values (transformed so that these values integrate to 1).

In principle, we could construct these estimates for all the parameters in the model. However, in practice the maximum likelihood estimates for most of the parameters in the model are well determined,<sup>5</sup> and it is only the correlation parameter  $\rho$  that poses problems. We therefore use a profile (or concentrated) likelihood for  $\rho$  (see Appendix 3 for details). The profile likelihood can be used to substitute for the full likelihood when there are other parameters in the model which are not of direct interest [14]. We have verified that estimates for these other parameters are very stable across all models. Therefore, once  $\rho$  is estimated we find the maximum likelihood values of the other parameters given this value of  $\rho$ .

We report the maximum likelihood estimates and our new bias correction estimates together with bootstrapped standard errors and confidence intervals. Our cluster bootstrapping takes account of the stratification and cluster sampling procedure of the survey design [15]. All our HIV prevalence estimates are weighted.

---

<sup>5</sup>Once we have accounted for the difficulties associated with the interviewer parameters using the random effects approach.

## A2 Additional Results for Model Parameters

We begin by estimating the interviewer random effects as the average error term (for each interviewer) from a probit model for consent to test for HIV, where we include a standard set of covariates along with the mean of these variables for each interviewer to capture the effects of non-random allocation of interviewers to participants as shown in equation (A4). Descriptive statistics for interviewers in Zambia and Ghana are shown in Table A1. Table A2 lists the variables used as predictors in the bivariate probit model, and their source in the data. The Stata code for preparing the data is publicly available from <http://hdl.handle.net/1902.1/17657> [3, 16]. The distribution of the number of interviewees by interviewer and their consent rates is shown in figures A1-A4.

Table A3 presents results from the probit model for consent to test for HIV in Zambia. This table is a single regression for HIV consent where the first column shows the marginal effects for the individual level variables ( $x_{ij}$ ), while the corresponding marginal effects for the interviewer averages ( $\bar{x}_j$ ) are shown in the second column.<sup>6</sup> At the individual level, the following variables are associated with consent to test: education, location, prior sexually transmitted disease (STDs), age at first intercourse, number of partners, willingness to care for a HIV positive relative, knowing an AIDS victim, and being a smoker. For example, an extra year of education is associated with an increase in the probability of consenting to an HIV test by roughly 0.5 percentage points. Apart from years of education, which is the mean years of education for that interviewer’s interviewees, all interviewer averages are measured as the proportion of interviewees in that category. Many of the interviewer average measures are significant in predicting consent. For example, in table A3 if the interviewee speaks Lozi the probability of consent to testing increases, but the association is not statistically significant. However, interviewers who conducted more interviews with Lozi speakers had higher consent rates (indicated by the positive and statistically significant coefficient for the proportion of interviews with Lozi speakers in table A3). These interviewers are likely to be Lozi speakers themselves, and these results therefore suggest that the Lozi speaking interviewers may have been better than average at obtaining consent.

Table A4 presents marginal effects for the Heckman-type selection model for consent and HIV status for Zambia. The first two columns give results for the maximum likelihood fixed effects approach. The middle two columns give results for our maximum likelihood random effects approach, and the final two columns give the results for the random effects bias correction model. For this third model, we estimate the likelihood on a grid of values of  $\rho$  (we use values between  $-1$  and  $+1$  at intervals of 0.01). We then calculate the likelihood of the model and posterior probability of each value of  $\rho$ , and then find the expected value of  $\rho$  based on this probability.<sup>7</sup> The coefficients reported in table A4 are calculated with this value of  $\rho$  imposed. Coefficient estimates across the three models are very similar.

Chiburis and colleagues recommend the use of the bootstrap for inference in the context of recursive bivariate probit models to correct for poor coverage of analytic standard errors [9]. We find that the bootstrap confidence interval for the random effects model is almost 10 times as wide as the analytic standard errors for the fixed effects model. We use 1,000 iterations to calculate

---

<sup>6</sup>Marginal effects show the expected change in the probability of a positive outcome for a unit change in that covariate, evaluated at the mean values of the covariates.

<sup>7</sup>The posterior probability is calculated by applying a constant to the likelihood for each value of  $\rho$  such that these transformed likelihoods integrate to 1.

these standard errors, and the corresponding 95% bootstrap confidence interval is calculated using the appropriate centiles from the empirical distribution of the bootstrap estimates. This approach is more appropriate than normal-based approximations when the distribution of the parameter of interest is skewed.

In table A5, we present the estimated correlation coefficient for Zambia from the different models. The negative values estimated indicate that those who refuse to test are more likely to be HIV-positive. The maximum likelihood random effects approach yields a  $\rho$  of around  $-0.50$ , which is somewhat lower than the  $-0.75$  obtained from the fixed effects estimator. The random effects bias corrected estimate is slightly smaller again at  $-0.44$ . Table A6 presents estimates of HIV prevalence among those who refused to consent to an HIV test, again comparing the fixed effect, random effect and random effects bias corrected methods. As with the correlation coefficient, the estimates from the random effects model are lower than from the fixed effects model (32% v 52%). The corresponding estimate from the imputation model is 12%. Table A7 presents marginal effects for the consent model in Ghana, and tables A8-A10 give the parameter results of the three Heckman-type selection models.

### A3 The concentrated (profile) likelihood function

The likelihood of the parameters  $(\beta_s, \beta_h, \rho)$ , given the full data set  $(y, x, z)$  is:

$$L(\beta_s, \beta_h, \rho) = P(y, x, z | \beta_s, \beta_h, \rho) \tag{A9}$$

For a given  $\rho$ , we can concentrate the likelihood function by setting the other parameters at their maximum likelihood values given  $\rho$ :

$$L_c(\rho) = P(y, x, z | \hat{\beta}_s(\rho), \hat{\beta}_h(\rho), \rho) \approx P(y, x, z | \rho) \tag{A10}$$

In large samples, the approximation to  $P(y, x, z | \rho)$  will become exact as the maximum likelihood estimates of the other parameters are consistent. Using the concentrated maximum likelihood, the problem is reduced to a one parameter model and we can carry out our small sample correction with a prior probability distribution over  $\rho$  alone.

## A4 Additional Tables and Figures

Table A1: Descriptive Statistics by Interviewer

	Number of Interviewers	Number of Interviews by Interviewer	Consent by Interviewer Median (Interquartile Range)	HIV Prevalence By Interviewer
Zambia	89	25 (4 – 132)	84% (72% – 93%)	10% (0% – 16%)
Ghana	55	114 (5 – 152)	87% (81% – 94%)	1% (0% – 2%)

Note to table A1: For each interviewer, their consent rate is calculated as the number of respondents from whom consent to test for HIV was obtained by the interviewer, divided by the number of respondents from whom consent to test for HIV was sought by the interviewer. For each interviewer, their HIV prevalence rate is calculated as the number of HIV positive respondents among those who consented to test for that interviewer, divided by the number of respondents who consented to test for that interviewer.

Table A2: DHS Variable Description

Variable	DHS Variable Name (Zambia 2007)	Note
<b>Demographic Characteristics</b>		
Age Category	v013	5 year age bands
Years of Education	v133	Education in single years
Marital Status	v501	Never/previously/currently married
Ethnicity	v131	Categorical
Language	smlangi	Language of interview
Religion	v130	Categorical
<b>Sex and Behavior</b>		
Would Respondent Care for HIV Positive Relative?	v778	Binary Indicator
Does Respondent Know Anyone Who Died of AIDS?	v775	Binary Indicator
Ever had an HIV Test	v781	Binary Indicator
Smoker	v463a, v463b, v463c	Binary Indicator
Alcohol Drinker	s1012a	Binary Indicator
Ever Had STD	v763a, v763b, v763c	Binary Indicator
Age at First Intercourse	v525	Never, <age 15, >age 15
Had High Risk Sex	v766a	Binary, sex with someone other than partner
Used Condom Last Intercourse	v761	Binary Indicator
Number of Partners Last 12 Months	v766a, v766b	Continuous
<b>Household Characteristics</b>		
Wealth Quintile	mv190	Categorical
Type of Location	hv025	Large City/Small City/Town/Countryside
Region	v101	Categorical
<b>HIV Testing</b>		
Consented to Test for HIV	ha63	Binary Indicator
Valid HIV Test Result	hb63	Binary Indicator
<b>Survey Characteristics</b>		
Interviewer Identity	v028	Anonymised ID Code
Primary Sampling Unit	v001	PSU Cluster
Household Weight	hv005	Nationally representative weight

Table A3: Marginal Effects from a Probit Model for HIV Test Consent for Zambian Men

Variable	HIV Test Consent	
	Individual Level Variables	Interviewer Level Averages
Years of Education	0.00475** (0.002)	-0.0001 (0.000)
<b>Wealth Category</b>		
Poorest	-0.00731 (0.031)	-0.0086*** (0.002)
Poorer	-0.02721 (0.031)	-0.0046 (0.003)
Middle	-0.03732 (0.027)	-0.0059 (0.004)
Richer	-0.00842 (0.019)	-0.0109*** (0.004)
<b>Location</b>		
Small city	0.12471*** (0.037)	0.0039 (0.004)
Town	0.14458*** (0.044)	0.0036 (0.004)
Countryside	0.20348*** (0.062)	0.0008 (0.003)
<b>Marital Status</b>		
Never Married	0.01165 (0.028)	0.0036 (0.003)
Currently Married	0.03447 (0.031)	-0.0018 (0.004)
Had STD	0.04105* (0.022)	0.0036 (0.003)
<b>Age at First Intercourse</b>		
15 or Younger	0.04986** (0.020)	0.0043 (0.003)
>15	0.03778* (0.022)	0.0013 (0.003)
<b>Had High Risk Sex</b>	0.05241** (0.025)	-0.0038 (0.004)
<b>Number of Partners</b>		
None	0.05095** (0.025)	-0.0078** (0.004)
2+	0.01516 (0.021)	0.0021 (0.003)
Used Condom Last Intercourse	-0.00033 (0.016)	-0.0005 (0.002)

Table A3 – Continued on the Next Page



HIV Test Consent		
Variable	Individual Level Variables	Interviewer Level Averages
Would Care for HIV Relative	0.05757** (0.024)	0.003 (0.003)
Knows Someone Who Died of AIDS	0.03795*** (0.011)	-0.0011 (0.001)
Previously HIV Tested	0.00504 (0.013)	-0.0038 (0.002)
Smoker	0.03644*** (0.012)	0.0041** (0.002)
Drinks Alcohol	0.01008 (0.012)	-0.0064*** (0.002)
<b>Language</b>		
English	-0.04601 (0.062)	0.0079 (0.008)
Bemba	-0.00834 (0.061)	0.0037 (0.009)
Lozi	0.04688 (0.056)	0.0130** (0.007)
Nyanja	0.05083 (0.055)	0.0036 (0.009)
Tonga	0.06172 (0.062)	0.006 (0.008)
Observations		6,416

Robust standard errors in parentheses

\*\*\* p < 0.01, \*\* p < 0.05, \* p < 0.1

Note to table A3: Age group, region, religion, and ethnicity are included in the regression but are not shown. Each column is for the same regression where HIV consent is regressed on the X variables and the corresponding average (in percent) for each interviewer. A probit model is used, and marginal effects (absolute change in the probability of a positive outcome) are shown. In some cases standard errors are very small, but not zero as appears in the table due to rounding. Interviewer averages of categorical variables are a proportion ranging from 0 to 1. Standard errors account for clustering at the PSU level.

Table A4: Marginal Effects for Heckman-type Selection Models (Men Zambia 2007)

Variables	Maximum Likelihood Fixed Effects		Maximum Likelihood Random Effects		Bias Correction Random Effects	
	Consent Equation	HIV Equation	Consent Equation	HIV Equation	Consent Equation	HIV Equation
Interviewer Effect				0.2845*** (0.028)		0.2857*** (0.028)
Years of Education	0.0046** (0.002)	0.0012 (0.002)	(0.002)	0.0019 (0.002)	0.0047** (0.002)	0.002 (0.002)
<b>Wealth Category</b>						
Poorest	-0.0172 (0.020)	0.0486** (0.020)	-0.0199 (0.020)	0.0463** (0.018)	-0.0198 (0.020)	0.0451** (0.018)
Poorer	-0.0262 (0.021)	0.0528*** (0.020)	-0.0299 (0.021)	0.0479*** (0.018)	-0.0298 (0.021)	0.0465*** (0.018)
Middle	0.0028 (0.025)	0.0645*** (0.024)	-0.0017 (0.024)	0.0630*** (0.022)	-0.0017 (0.024)	0.0620*** (0.021)
Richer	0.0055 (0.031)	0.0503* (0.028)	0.0072 (0.030)	0.0493* (0.025)	0.0071 (0.030)	0.0485* (0.025)
<b>Location</b>						
Small City	0.1529** (0.067)	-0.0589 (0.045)	0.1626** (0.064)	-0.0351 (0.035)	0.1622** (0.064)	-0.0308 (0.032)
Town	0.1605*** (0.061)	-0.0764** (0.038)	0.1649*** (0.059)	-0.0495* (0.027)	0.1647*** (0.059)	-0.0449* (0.024)
Countryside	0.1911*** (0.063)	-0.1289*** (0.039)	0.1972*** (0.059)	-0.0956*** (0.028)	0.1967*** (0.059)	-0.0896*** (0.024)
<b>Marital Status</b>						
Currently Married	0.0105 (0.029)	0.0787*** (0.030)	0.022 (0.028)	0.0791*** (0.027)	0.0221 (0.028)	0.0779*** (0.027)
Formerly Married	-0.0146 (0.028)	0.2125*** (0.026)	-0.0093 (0.028)	0.1992*** (0.025)	-0.0096 (0.028)	0.1950*** (0.023)
Had STD	0.0407 (0.025)	0.1079*** (0.023)	0.0437* (0.025)	0.1073*** (0.019)	0.0437* (0.025)	0.1064*** (0.019)
<b>Age at First Intercourse</b>						
15 or Younger	0.0502** (0.022)	-0.0422 (0.035)	0.0485** (0.021)	-0.0254 (0.031)	0.0487** (0.021)	-0.0228 (0.030)
>15	0.0374* (0.022)	-0.0475 (0.034)	0.0356* (0.021)	-0.0334 (0.030)	0.0357* (0.021)	-0.0312 (0.030)
<b>Had High Risk Sex</b>	0.0531* (0.027)	-0.0275 (0.025)	0.0546** (0.027)	-0.0199 (0.022)	0.0544** (0.027)	-0.0184 (0.021)
<b>Number of Partners</b>						
1	-0.0433 (0.027)	-0.0202 (0.029)	-0.0520* (0.027)	-0.0266 (0.026)	-0.0521* (0.027)	-0.0272 (0.025)
2+	-0.0259 (0.042)	0.0343 (0.039)	-0.0371 (0.042)	0.0293 (0.034)	-0.0371 (0.042)	0.0282 (0.033)
Condom Last Intercourse	-0.0079 (0.016)	0.0629*** (0.015)	-0.0017 (0.015)	0.0586*** (0.014)	-0.0014 (0.015)	0.0575*** (0.013)
Would Care for HIV Relative	0.0487** (0.020)	0.017 (0.033)	0.0524*** (0.020)	0.0285 (0.030)	0.0525*** (0.020)	0.0298 (0.029)
Know Someone Who Died of AIDS	0.0369*** (0.010)	-0.0144 (0.014)	0.0370*** (0.010)	-0.0081 (0.011)	0.0370*** (0.010)	-0.007 (0.011)
Previously HIV Tested	0.0053 (0.013)	0.0340*** (0.013)	0.005 (0.013)	0.0333*** (0.011)	0.005 (0.013)	0.0328*** (0.011)
Smoker	0.0371*** (0.013)	-0.0262* (0.016)	0.0371*** (0.013)	-0.0197 (0.013)	0.0370*** (0.013)	-0.0183 (0.013)
Drinks Alcohol	0.0097 (0.012)	0.0186 (0.012)	0.0106 (0.012)	0.0184* (0.011)	0.0105 (0.012)	0.0182* (0.011)
<b>Language</b>						
English	0.0276 (0.028)	0.0028 (0.033)	0.035 (0.029)	0.0075 (0.029)	0.0351 (0.029)	0.0081 (0.029)
Bemba	0.0953 (0.061)	0.0666 (0.049)	0.0916 (0.058)	0.0816* (0.042)	0.0919 (0.058)	0.0826** (0.041)
Lozi	0.0980*** (0.030)	-0.0076 (0.031)	0.0967*** (0.028)	0.0079 (0.026)	0.0965*** (0.028)	0.0099 (0.025)
Nyanja	0.1017** (0.046)	-0.0177 (0.043)	0.1133*** (0.036)	0.0035 (0.035)	0.1132*** (0.036)	0.0066 (0.034)
Tonga	0.0376 (0.048)	-0.0119 (0.064)	0.0453 (0.051)	-0.0132 (0.063)	0.0452 (0.051)	-0.0128 (0.062)
Observations	6,416	6,416	6,416	6,416	6,416	6,416

Robust standard errors in parentheses  
\*\*\* p < 0.01, \*\* p < 0.05, \* p < 0.1

Note to table A4: Controls for age group, region, ethnicity, and religion are included in each regression equation but not shown. The model is a bivariate probit for HIV status and consent to a HIV test. Data are for men who completed an interview. The first model uses interviewer fixed effects as the exclusion restriction. The coefficients from the interviewer effects are not shown in the table. The second model uses the random effects procedure where HIV consent is regressed on the X variables, along with the mean values for each interviewer (see table A3). The average error term for each interviewer is added to the predicted value of the interviewer means, and used as the exclusion restriction in the HIV regression. The final model is the random effects bias correction procedure using the same exclusion restriction. Coefficients are obtained by restricting the value of  $\rho$  to its random effects bias correction estimate and implementing the bivariate probit at that value. Marginal effects (absolute change in the probability of a positive outcome) are shown. Source: DHS Zambia 2007 (men). Standard errors account for clustering at the PSU level.

Table A5: Correlation Coefficient ( $\rho$ ) Estimates for Men in Zambia

Model for Refused Consent	Parameter Value	Analytic 95% CI		Bootstrap 95% CI	
Fixed Effects Model	-0.75	-0.94	-0.22		
Random Effects Model	-0.50	-0.73	-0.18	-0.68	0.23
Random Effects Bias Correction Model	-0.44			-0.64	0.43

Note to table A5: The table shows the estimated correlation coefficient between consent and HIV status for the fixed effects, random effects and random effects bias correction models. Analytic standard errors are shown for the fixed effects and random effects models, with bootstrapped errors for random effects and random effects bias correction models using 1,000 replications. Source: DHS Zambia 2007 (men).

Table A6: HIV Prevalence (%) among Men who Refused to Test in Zambia

Model for Refused Consent	Parameter Value	Analytic 95% CI		Bootstrap 95% CI	
Fixed Effects Model	52.00%	49.90%	54.00%		
Random Effects Model	32.00%	30.30%	33.80%	6.30%	42.70%
Random Effects Bias Correction Model	28.60%	26.60%	29.80%	3.20%	39.80%
Imputation Model	11.70%	10.80%	12.60%		

Note to table A6: The table shows the estimated HIV prevalence rate among individuals who refused consent for the fixed effects, random effects and random effects bias correction models. Analytic standard errors are shown for the fixed effects and random effects models, with bootstrapped errors for random effects and random effects bias correction models using 1,000 replications. Also shown is the HIV rate using an imputation model for men who refused consent. Prevalence estimates are weighted and account for survey design. Source: DHS Zambia 2007 (men).

Table A7: Marginal Effects from a Probit Model for HIV Test Consent for Ghanaian Men

HIV Test Consent		
Variables	Individual Level Variables	Interviewer Level Averages
Years of Education	0.00224* (0.001)	-0.0002 (0.000)
<b>Wealth Category</b>		
Poorest	0.09640*** (0.016)	0.0136** (0.007)
Poorer	0.08063*** (0.014)	0.0200*** (0.001)
Middle	0.05346*** (0.015)	0.0417*** (0.007)
Richer	0.03687*** (0.014)	0.0376*** (0.006)
<b>Marital Status</b>		
Currently Married	0.02424 (0.022)	-0.0035 (0.004)
Formerly Married	-0.04371 (0.028)	0.0084 (0.011)
Had STD	0.03830* (0.021)	0.0529*** (0.009)
<b>Age at First Intercourse</b>		
Never Had Sex	0.00918 (0.021)	-0.0578*** (0.009)
15 or Younger	-0.02929 (0.018)	-0.0569*** (0.010)
<b>Had High Risk Sex</b>		
	0.00272 (0.019)	0.0092 (0.007)
<b>Number of Partners</b>		
1	0.00304 (0.022)	-0.0331*** (0.005)
2+	-0.00752 (0.031)	0.0218** (0.010)
Condom Last Intercourse	0.00045 (0.017)	-0.1715*** (0.024)
Would Care for HIV Relative	0.02213* (0.013)	0.0124*** (0.001)
Know Someone Died of AIDS	0.01538 (0.010)	-0.0424*** (0.006)
Previously HIV Tested	0.00564	-0.0097***

Table A7 – Continued on the Next Page

---

HIV Test Consent

Variables	Individual Level Variables	Interviewer Level Averages
	(0.017)	(0.003)
Smoker	-0.00928	-0.0233***
	(0.017)	(0.007)
<b>Language</b>		
Akan	0.04955***	-0.0147***
	(0.017)	(0.001)
Ga	0.06397**	0.0282***
	(0.026)	(0.005)
Ewe	0.02522	0.0175***
	(0.038)	(0.005)
Nzema	0.043	0.0112*
	(0.047)	(0.007)
Dagbani	0.01586	-0.0241***
	(0.049)	(0.007)
Other	-0.03509	
	(0.030)	
Observations	4,955	

Robust standard errors in parentheses

\*\*\* p < 0.01, \*\* p < 0.05, \* p < 0.1

Note to table A7: Age group, region, religion, and ethnicity are included in the regression but are not shown. Each column is for the same regression where HIV consent is regressed on the X variables and the corresponding average (in percent) for each interviewer. A probit model is used, and marginal effects (absolute change in the probability of a positive outcome) are shown. In some cases standard errors are very small, but not zero as appears in the table due to rounding. Interviewer averages of categorical variables are a proportion ranging from 0 to 1. Standard errors account for clustering at the PSU level.

Table A8: Marginal Effects for Heckman-type Selection Models (Ghana 2003)

Variables	Maximum Likelihood Fixed Effects		Maximum Likelihood Random Effects		Bias Correction Random Effects	
	Consent Equation	HIV Equation	Consent Equation	HIV Equation	Consent Equation	HIV Equation
Interviewer Effect				0.2125*** (0.020)		0.2120*** (0.020)
Years of Education	0.002 (0.001)	0.0004 (0.000)	0.0023* (0.001)	0.0004 (0.000)	0.0023* (0.001)	0.0004 (0.000)
<b>Wealth Category</b>						
Poorest	-0.0187 (0.020)	0.0039 (0.005)	-0.0184 (0.019)	0.0039 (0.005)	-0.0184 (0.019)	0.004 (0.005)
Poorer	-0.0563*** (0.020)	0.0078 (0.006)	-0.0553*** (0.020)	0.0078 (0.006)	-0.0554*** (0.020)	0.0079 (0.006)
Middle	-0.0747*** (0.021)	0.0017 (0.006)	-0.0769*** (0.021)	0.0017 (0.006)	-0.0769*** (0.021)	0.0018 (0.006)
Richer	-0.1139*** (0.024)	0.0021 (0.007)	-0.1169*** (0.023)	0.0021 (0.007)	-0.1171*** (0.023)	0.0023 (0.007)
<b>Marital Status</b>						
Currently Married	0.0241 (0.022)	0.0128 (0.008)	0.0235 (0.022)	0.0128 (0.008)	0.0237 (0.022)	0.0128 (0.008)
Formerly Married	-0.0382 (0.024)	0.0088 (0.008)	-0.0420* (0.024)	0.0087 (0.008)	-0.0418* (0.024)	0.0088 (0.008)
Had STD	0.0461 (0.029)	0.0093 (0.007)	0.0457 (0.028)	0.0094 (0.007)	0.0456 (0.028)	0.0094 (0.007)
<b>Age at First Intercourse</b>						
15 or Younger	0.0502** (0.022)	-0.0422 (0.035)	0.0485** (0.021)	-0.0254 (0.031)	0.0487** (0.021)	-0.0228 (0.030)
>15	0.0374* (0.022)	-0.0475 (0.034)	0.0356* (0.021)	-0.0334 (0.030)	0.0357* (0.021)	-0.0312 (0.030)
<b>Had High Risk Sex</b>	0.0531* (0.027)	-0.0275 (0.025)	0.0546** (0.027)	-0.0199 (0.022)	0.0544** (0.027)	-0.0184 (0.021)
<b>Number of Partners</b>						
1	0.004 (0.022)	-0.01 (0.007)	0.0033 (0.022)	-0.01 (0.007)	0.0032 (0.022)	-0.0101 (0.007)
2+	-0.0064 (0.031)	-0.0178* (0.011)	-0.0068 (0.030)	-0.0180* (0.011)	-0.0069 (0.031)	-0.0180* (0.011)
Condom Last Intercourse	0.0001 (0.018)	0.0083* (0.004)	0.0000 (0.018)	0.0082* (0.004)	0.0000 (0.018)	0.0083* (0.004)
Would Care for HIV Relative	0.0221* (0.013)	-0.0007 (0.004)	0.0230* (0.012)	-0.0007 (0.004)	0.0229* (0.012)	-0.0007 (0.004)
Know Someone Who Died of AIDS	0.0163 (0.011)	-0.0009 (0.003)	0.0162 (0.011)	-0.001 (0.003)	0.0161 (0.011)	-0.0009 (0.003)
Previously HIV Tested	0.004 (0.018)	-0.0054 (0.007)	0.0046 (0.018)	-0.0054 (0.007)	0.0049 (0.018)	-0.0056 (0.007)
Smoker	-0.0103 (0.017)	0.0057 (0.004)	-0.0089 (0.017)	0.0057 (0.004)	-0.0089 (0.017)	0.0058 (0.004)
<b>Language</b>						
Akan	0.0474*** (0.017)	0.0042 (0.004)	0.0513*** (0.016)	0.0043 (0.004)	0.0511*** (0.016)	0.0044 (0.004)
Ga	0.0766* (0.045)	0.0009 (0.013)	0.0864** (0.044)	0.001 (0.013)	0.0858** (0.044)	0.001 (0.013)
Ewe	0.0338 (0.045)	0.0122 (0.009)	0.0272 (0.045)	0.0122 (0.009)	0.0271 (0.045)	0.0123 (0.009)
Nzema	0.0537 (0.060)	0.0253* (0.015)	0.0484 (0.052)	0.0253* (0.015)	0.0489 (0.052)	0.0254* (0.015)
Dagbani	-0.0201 (0.063)	-0.1142*** (0.018)	0.0187 (0.038)	-0.1387*** (0.021)	0.0186 (0.038)	-0.8298*** (0.084)
Other	-0.0391 (0.027)	-0.0102 (0.009)	-0.0327 (0.025)	-0.0102 (0.009)	-0.0328 (0.025)	-0.0102 (0.009)
Observations	4,955	4,955	4,955	4,955	4,955	4,955

Robust standard errors in parentheses  
 \*\*\* p < 0.01, \*\* p < 0.05, \* p < 0.1

Note to table A8: Controls for age group, region, ethnicity, and religion are included in each regression equation but not shown. The model is a bivariate probit for HIV status and consent to a HIV test. Data are for men who completed an interview. The first model uses interviewer fixed effects as the exclusion restriction. The coefficients from the interviewer effects are not shown in the table. The second model uses the random effects procedure where HIV consent is regressed on the X variables, along with the mean values for each interviewer (see table A3). The average error term for each interviewer is added to the predicted value of the interviewer means, and used as the exclusion restriction in the HIV regression. The final model is the random bias correction procedure using the same exclusion restriction. Coefficients are obtained by restricting the value of  $\rho$  to its random effects bias correction estimate and implementing the bivariate probit at that value. Marginal effects (absolute change in the probability of a positive outcome) are shown. Source: DHS Ghana 2003 (men). Standard errors account for clustering at the PSU level.

Table A9: Correlation Coefficient ( $\rho$ ) Estimates for Men in Ghana

Model for Refused Consent	Parameter Value	Analytic 95% CI		Bootstrap 95% CI	
Fixed Effects Model	0.93	0.50	0.99		
Random Effects Model	0.93	0.02	1.00	0.56	1.00
Random Effects Bias Correction Model	0.59			0.39	0.72

Note to table A9: The table shows the estimated correlation coefficient between consent and HIV status for the fixed effects, random effects and random effects bias correction models. Analytic standard errors are shown for the fixed effects and random effects models, with bootstrapped errors for random effects and random effects bias correction models using 1,000 replications. Source: DHS Ghana 2003 (men).


Table A10: HIV Prevalence (%) among Men who Refused to Test in Ghana

Model for Refused Consent	Parameter Value	Analytic 95% CI		Bootstrap 95% CI	
Fixed Effects Model	1.00E-07	3.00E-08	2.00E-07		
Random Effects Model	2.00E-05	1.00E-06	4.00E-06	1.00E-08	0.16%
Random Effects Bias Correction Model	0.07%	0.05%	0.09%	0.03%	0.35%
Imputation Model	1.82%	1.60%	20.48%		

Note to table A10: The table shows the estimated HIV prevalence rate among individuals who refused consent for the fixed effects, random effects and bias correction models. Analytic standard errors are shown for the fixed effects and random effects models, with bootstrapped errors for random effects and bias correction models using 1,000 replications. Also shown is the HIV rate using an imputation model for men who refused consent. Estimates lower than 0.01% are shown in scientific notation. Prevalence estimates are weighted and account for survey design. Source: DHS Ghana 2003 (men).

## Additional Figures


Figure A1: Histogram of Number of Interviews by Interviewer in Zambia 2007 (Men)



Graph is at the interviewer level (one observation per interviewer).




Figure A2: Histogram of Consent Rates by Interviewer in Zambia 2007 (Men)




Note to Figure A2: Graph is at the interviewer level (one observation per interviewer). For each interviewer, their consent rate is calculated as the number of respondents from whom consent to test for HIV was obtained by the interviewer, divided by the number of respondents from whom consent to test for HIV was sought by the interviewer.

Figure A3: Histogram of Number of Interviews by Interviewer in Ghana 2003 (Men)



Graph is at the interviewer level (one observation per interviewer).

Figure A4: Histogram of Consent Rates by Interviewer in Ghana 2003 (Men)



Note to Figure A4: Graph is at the interviewer level (one observation per interviewer). For each interviewer, their consent rate is calculated as the number of respondents from whom consent to test for HIV was obtained by the interviewer, divided by the number of respondents from whom consent to test for HIV was sought by the interviewer.

## References

1. W. P. Van de Ven and B. Van Praag. The demand for deductibles in private health insurance: A probit model with sample selection. *J Econom*, 17:229–252, 1981.
2. J. A. Dubin and D. Rivers. Selection bias in linear regression, logit and probit models. *Sociol Methods Res*, 18:360–390, 1989.
3. T. Bärnighausen, J. Bor, S. Wandira-Kazibwe, and D. Canning. Correcting HIV prevalence estimates for survey nonparticipation using Heckman-type selection models. *Epidemiology*, 22:27–35, 2011.
4. D. Madden. Sample selection versus two-part models revisited: the case of female smoking and drinking. *J. Health Econ.*, 27:300–307, 2008.
5. Y. Mundlak. On the pooling of time series and cross section data. *Econometrica*, 46:69–85, 1978.
6. G. Chamberlain. Analysis of covariance with qualitative data. *Rev Econ Stud*, 47:225–238, 1980.
7. K. M. Murphy and R. H. Topel. Estimation and inference in two-step econometric models. *J Bus Econ Stat*, 20:88–97, 2002.
8. J. S. Butler. Estimating the correlation in censored probit models. *Rev Econ Stat*, 78:356–358, 1996.
9. R. C. Chiburis, J. Das, and M. Lokshin. A practical comparison of the bivariate probit and linear IV estimators. *Econ Lett*, 117:762–766, 2012.
10. D. W. Andrews. Estimation when a parameter is on a boundary. *Econometrica*, 67:1341–1383, 1999.
11. D. W. Andrews. Inconsistency of the bootstrap when a parameter is on the boundary of the parameter space. *Econometrica*, 68:399–405, 2000.
12. W. Greene. The behaviour of the maximum likelihood estimator of limited dependent variable models in the presence of fixed effects. *Econom J*, 7:98–119, 2004.
13. S. Greenland. Small-sample bias and corrections for conditional maximum-likelihood odds-ratio estimators. *Biostatistics*, 1:113–122, 2000.
14. S. A. Murphy and A. W. Van der Vaart. On profile likelihood. *J. Am. Stat. Assoc.*, 95:449–465, 2000.
15. D. J. Corsi, M. Neuman, J. E. Finlay, and S. Subramanian. Demographic and health surveys: a profile. *Int. J. Epidemiol.*, 41:1602–1613, 2012.
16. D. R. Hogan, J. A. Salomon, D. Canning, J. K. Hammitt, A. M. Zaslavsky, and T. Bärnighausen. National HIV prevalence estimates for sub-saharan Africa: controlling selection bias with Heckman-type selection models. *Sex. Transm. Infect.*, 88:i17–i23, 2012.